

Agent-based modelling of stock markets using existing order book data

Efstathios Panayi¹, Mark Harman¹, Anne Wetherilt²

¹ UCL, Gower Street, London WC1E 6BT, UK,

² Bank of England, Threadneedle Street, London EC2R 8AH, UK**

Abstract. We propose a new method for creating alternative scenarios for the evolution of a financial time series over short time periods. Using real order book data from the Chi-X exchange, along with a number of agents to interact with that data, we create a semi-synthetic time series of stock prices. We investigate the impact of using both simple, limited intelligence traders, along with a more realistic set of traders. We also test two different hypotheses about how real participants in the market would modify their orders in the alternative scenario created by the model. We run our experiments on 3 different stocks, evaluating a number of financial metrics for intra- and inter-day variability. Our results using realistic traders and relative pricing of real orders were found to outperform other approaches.

1 Introduction

The behaviour of stock prices over short horizons is an important consideration for both market participants and regulators, as the former need to be confident in their ability to place and execute their orders, while the latter need to ensure a smoothly functioning market. Currently, estimating the range of prices that could arise in the short term is predominantly focused on analysing past data and fitting statistical models to specific time series from which they come. This approach is based on the assumption that past market conditions are likely to be repeated at some point in the future. However, it is very restrictive, as it does not allow for scenarios that have not previously occurred [20], or have occurred only very rarely (such as large, rapid intra-day movements). Unfortunately, this restriction is highly problematic, because such infrequent, rapid intra-day movements denote one of the largest market risks.

Agent-based modelling is a well established method for creating alternative scenarios in a financial market, the first work on this being conducted 3 decades ago [5]. The agent-based approach seeks to program the behaviour of individual traders, and their interaction gives rise to changes in the intra-day behaviour of orders and prices. Agent-based modelling offers many more parameters that can

** The views expressed in this paper are the authors' and are not necessarily those of the Bank of England.

be altered to generate different scenarios. It also facilitates the study of emergent properties of traders' interactions and particular classes of traders in isolation.

Agent-based modelling has, so far, been used primarily for stock market simulation with a focus on longer time frames than intra-day behaviour. A key challenge for agent-based models is to demonstrate that the resulting price dynamics are indeed consistent with known empirical facts. Such 'stylised facts' may include volatility clustering and 'fat tails' in distributions of financial returns (Engle and Russell [18]).

In this paper, we seek to combine realistic data and an agent-based model to achieve a simulation that exploits real world data. We start with 3 high quality data sets from the Chi-X exchange and rebuild the order book so that we can pause the market at any time and examine the bids and offers for the stock, along with any order executions and cancellations. We then add different classes of agents to interact with this 'live' order book, so that the evolution of the stock price is modified by the interaction with our agents.

In particular, in this paper we introduce and empirically compare a class of almost zero-intelligence traders, along with a class of traders that is based on more realistic behaviour and compare their impact on the stock price. We also experiment and report on two different ways of incorporating real data into the model: absolute and relative pricing. The former assumes that if traders in the real market were participating in the synthetic market, they would have submitted their orders at exactly the same prices they had originally. The latter method assumes that traders would have shifted the price of their orders by the difference between the stock price in the synthetic market and in the real market.

We empirically evaluate the behaviour of the stock price resulting from the model for 3 of the most frequently traded stocks on the Chi-X exchange (Arcelor Mittal, Deutsche Bank and GDF Suez). We do this by comparing the ranges of maximum, minimum and closing prices produced by multiple runs of the model, and running tests for fat tails and volatility clustering of the returns distribution.

This paper contributes to the existing literature on agent-based modelling of financial markets in two ways: Firstly, it introduces the concept of semi-synthetic modelling, which combines past data and agents in a single simulation, and intuitively should be closer to the real market than a pure agent-based model. Secondly, whereas existing studies have longer horizons, we focus on short-term behaviour of stock prices and study the returns at the transaction level. Our interest in short-term market behaviour is motivated by the rapid intra-day drop and recovery in the US equities market, during the May 6, 2010 'Flash Crash', which created concerns about short term proprietary trading behaviour [9].

The rest of this paper is organised as follows: Section 2 introduces the experimental framework employed in this study and describes the two classes of traders that we compare, along with the two ways of handling the real order book data. Section 3 presents the research questions and gives some detail about the stocks we study. Section 4 then presents the results of batch runs of our model. Section 5 suggests potential weaknesses of our study. Section 6 summarises related work in the area of agent-based modelling of markets. Section 7 concludes.

2 Agent-Based Simulation with ORder Book Data (ABSORBD)

Our model aims to replicate the activity of a single day on the Chi-X exchange. The two main types of orders are market orders, where a trader can buy or sell a particular amount of stock at the best price available at the moment, or a limit order, where the trader specifies a price above which she is unwilling to buy, if she submits a bid, or a price under which she is unwilling to sell, if she submits an offer for a stock. If a limit order isn't executed immediately in its entirety, it enters the order book, where bids and offers are prioritised by price, then by time, in the case of tie breaks.

2.1 Stylised facts

In agent-based models of financial markets, it is standard practice to measure the validity of the model by investigating whether the stock price exhibits particular characteristics, known as the 'stylised facts'. There are a number of ways to replicate these characteristics, a summary of which is presented in Section 6.

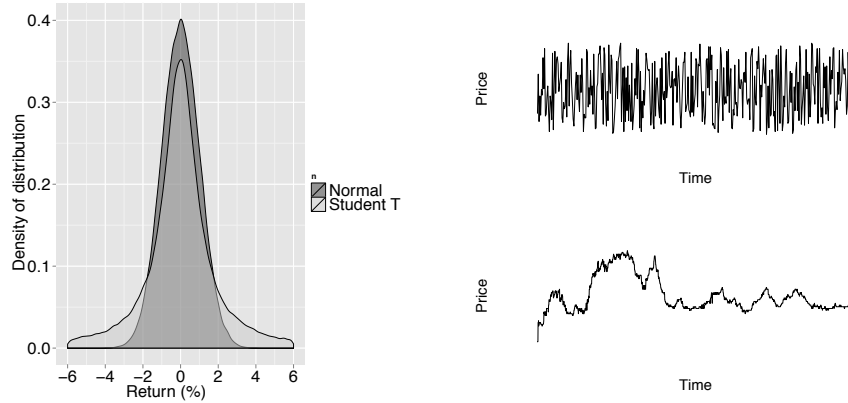
While the literature about stylised facts, which commenced with Mandelbrot [14], was initially concerned with characteristics of markets at longer time scales, Engle and Russell [18] demonstrated that these characteristics also apply to the intra-day level. This paper studies the behaviour of prices at a very high frequency, namely the change in price between two subsequent transactions. We evaluate our model based on the following two widely used stylised facts, also illustrated in Figure 1:

1. Fat tails. This means that, when plotting the distribution of returns of the financial asset, the probability of very high or very low returns is higher than that implied by a normal distribution with the same mean and standard deviation. The degree of 'fat-tailedness' is called kurtosis, and is the fourth moment of the distribution. The normal distribution has a kurtosis of 3, and a distribution with a kurtosis above 3 is called leptokurtic, or fat-tailed.
2. Volatility clustering. This means that a large change in the asset price (over a minute, for example) is more likely to be followed by a large change, and the same is true for small changes in the asset price. The time series for which this is true are called 'heteroscedastic' (i.e. of differing variability), and we can test whether a series is heteroscedastic with the Engle ARCH-LM test.

2.2 Synthetic trader behaviour

In our model, we create alternative scenarios for the evolution of a stock price over a day by recreating the order book for a single stock with real orders and then adding synthetic traders who also submit orders. We investigate the impact of agent behaviour on the stock price and other market characteristics, described in 2.1. In particular, we implement 2 different classes of synthetic traders and compare their impact on the market.

IV



Normal and Student T distributions. Note the shorter middle and fatter tails of the Student T distribution.

Homoscedastic and heteroscedastic time series. The variability is constant in the upper (homoscedastic) time series, but changes in the lower (heteroscedastic) plot.

Fig. 1: Fat-tails and volatility clustering, two of the most common characteristics of financial time series

Our first class of traders consists of homogeneous ‘Limited Intelligence’ (LI) traders, similar to those proposed by LiCalizi and Pelizzari [12] and in the physics literature, for example by Maslov [15]. LI traders’ decisions are unaffected by the stock price and only take into account their budget constraints when determining whether to submit a trade. We call these limited, as opposed to zero intelligence traders, as they don’t place orders at prices that are worse than are available in the prevailing market and they stop when they have reached their cash limit. Although their trading behaviour would not arise from any meaningful strategy, is useful to implement as a baseline against which to compare it to other behaviours.

In our model, LI traders become active after a certain amount of trading activity in the market, which varies for each trader. When a particular trader becomes active, she decides whether to submit a bid or an offer for the stock, with $X\%$ and $(100-X)\%$ probability respectively. If she is still within her budget, she continues with her order and determines the price and size randomly: The order price is uniformly distributed in the range of $Y\%$ below or above the best bid or offer respectively, while the size (in shares) is uniformly distributed within the limits $[1, Z]$.

Our second class of traders is based on more realistic trader behaviour. Kirilenko et al. [9] study the composition of the E-mini S&P 500 stock index futures market, and identify 6 major categories of market participants:

1. Intermediaries, who usually post prices on both sides of the order book and try to maintain their position throughout the day, making their income from the difference between their bid and offer price.
2. High Frequency Traders, who have a relatively low net position throughout the day, compared to their activity. They are similar to intermediaries, but have much higher trading activity and much shorter holding periods.
3. Fundamental Buyers, who try to build a long position during the day.
4. Fundamental Sellers, who try to build a short position during the day.
5. Opportunistic Traders, who may behave as intermediaries at times, or as fundamental traders at times when they see significant directional moves.
6. Small Traders, who show very limited trading activity.

We implement all of these categories (except the Small Traders, who have very little, if any, effect on the stock price) in our second class and will refer to the traders as Kirilenko (KI) traders collectively. Our first class comprises 1000 LI traders. Our second class consists of 394 KI traders, distributed as shown in in Table 1. These numbers were chosen based on the classification in [9] and to produce the empirical features in 2.1.

Trader type	Number
Intermediary	40
HFT	4
Opportunistic	150
Fundamental Buyer	100
Fundamental Seller	100

Table 1: Trader Information

2.3 Handling of real order book data

As we are adding synthetic traders to interact with our full order book data from the Chi-X exchange, we have to make some assumptions about how the real traders would have reacted to the modified stock prices. As the order book only provides anonymous trading data, we cannot identify individual trading strategies, and hence we need to make assumptions as to how they would have interacted with either LI or KI traders.

To do so, we consider two approaches. For our first method (referred to as the absolute pricing method), we assume that if the ‘real’ traders (i.e. the traders that submitted the orders in our dataset) were participants in our synthetic market, they would have submitted their orders at exactly the same price, independently of what happened to the stock price during the day. The second method (labelled relative pricing method) assumes that traders submit orders with prices that are relative to the prevailing stock price. So if the price in the synthetic stock market is higher, the orders read in from the real dataset will also be higher by the same relative amount. This is perhaps a more realistic assumption about how market participants would react to seeing different stock prices.

3 Experimental setup

As we are creating a semi-synthetic model, it seems reasonable to aim for a market where approximately half of the trades come from the real dataset, with the other half coming from our synthetic market participants. Since this precise split cannot be achieved in every run, we allow for runs where the number of trades in which at least one of the traders is synthetic is at least 30%. We adjust the number of traders and frequency with which they visit the market, in order to achieve this split.

3.1 Dataset description

This paper uses full order book data for 3 stocks from the Chi-X exchange on 3/1/2011, for which some detail is presented in Figure 2. These datasets are very detailed and contain every order to buy or sell a stock that was submitted to the exchange, the size of those orders, the time of submission and any executions or cancellations.

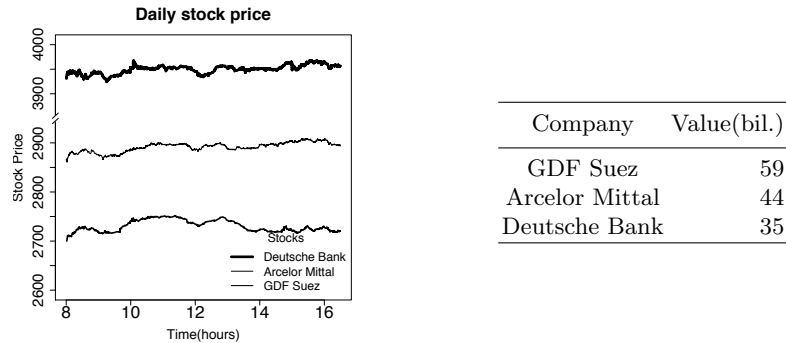


Fig. 2: The evolution of the stock prices for the 3 companies studied (on 3/1/2011) and company information

3.2 Research questions

In order to study the validity of our semi-synthetic model, we measure the impact on 5 inter- and intra-day variability measures:

- Closing, maximum and minimum prices observed over the day
- Kurtosis and volatility clustering.

These metrics are then used in our research questions, in order to investigate the impact of our design choices. These are as follows:

1. What is the effect on the market metrics mentioned above, of using absolute and relative pricing?
2. What is the effect on the market metrics mentioned above, of using simple (LI) and realistic (KI) traders?

In particular, for the first 3 metrics, we use a separate dataset of 60 points of coarse-grained (daily) data around our trading day from Yahoo! Finance, in order to construct boxplots of the difference between the closing, maximum and minimum prices and the starting price of the day. We then produce similar boxplots from 60 simulation runs for each design choice and compare the results of each batch run.

In more detail, the summary output of the coarse-grained, real dataset is compared to the summary outputs of separate batch runs where we use:

1. Absolute pricing and LI traders
2. Relative pricing and LI traders
3. Absolute pricing and KI traders
4. Relative pricing and KI traders

To measure volatility clustering, we use the Engle ARCH-LM test for conditional heteroscedasticity, whose null hypothesis is that a particular return distribution exhibits no ARCH effects (i.e. volatility is constant throughout the day). We study which of the four options mentioned above will produce runs in which this null hypothesis is rejected (and thus, there is clustering of volatility).

4 Results and discussion

As space does not allow us to include all the results of our simulation, Table 2 and the boxplots in Figure 4 give us summary information, while Figure 3 demonstrates four typical runs, which allow us to zoom in on the data. We see that in the second graph of Figure 3, we have an almost constant upward trend and we find that all runs that use absolute pricing and LI traders result in either upward or downward trends throughout the day. This is because of the homogeneous construction of the traders, and leads to stock price moves that are similar in nature (but not in amplitude). This also means that the volatility is more constant throughout the day, compared to the other methods.

Figure 4 shows the variability of the runs, in terms of the difference between the starting and closing, minimum and maximum prices of the day. We see that using absolute pricing greatly constrains the range of maximum, minimum and closing prices we observe in our runs, whereas it seems that using either LI or KI traders and relative pricing produces results that are reasonably close to the range of real prices. We see that our batch runs that use KI traders, however, produce daily maxima that are significantly higher than that we observe in the real dataset. So the combination of relative pricing and KI traders must give rise to higher intra-day volatility, something which we confirm by looking at the individual runs.

Next, investigate the two market characteristics often observed in the real markets, namely fat tails and volatility clustering. The distribution of returns from our model has fatter tails than the normal distribution, as the kurtosis is much higher than 3, using either LI or KI traders, and under either absolute or




Fig. 3: Four sample stock runs, using absolute or relative pricing and LI traders(top) and absolute or relative pricing and KI traders(bottom)

relative pricing. The relative order book pricing/KI traders combination, however, produces time series with higher kurtosis, as we can see in Table 2. We have already noted that this combination produces a more volatile stock price, and the aggressive nature of some of the traders also causes price jumps, and thus higher kurtosis.

	GSZ _p	MT _a	DBK _d
Real Order book data	29.0	10.1	8.8
with LI traders (abs)	45.2	13.0	35.1
with LI traders (rel)	43.5	13.0	43.5
with KI traders (abs)	36.9	15.3	32.3
with KI traders (rel)	81.3	152	99.7

Table 2: Summary kurtosis information

We also test for volatility clustering, using the ARCH-LM test proposed by Engle [7]. For 3 of the 4 options, with high probability, the test rejects the hypothesis that the time series of returns exhibits no ARCH effects. Therefore, the time series must exhibit differing volatility at various points in the trading day, which is what we observe in real markets also. For the combination of relative pricing and LI traders, however, we have some runs in which the hypothesis is not rejected, which means there is no sufficient evidence that volatility varies throughout the day. We have already mentioned that this result comes from the

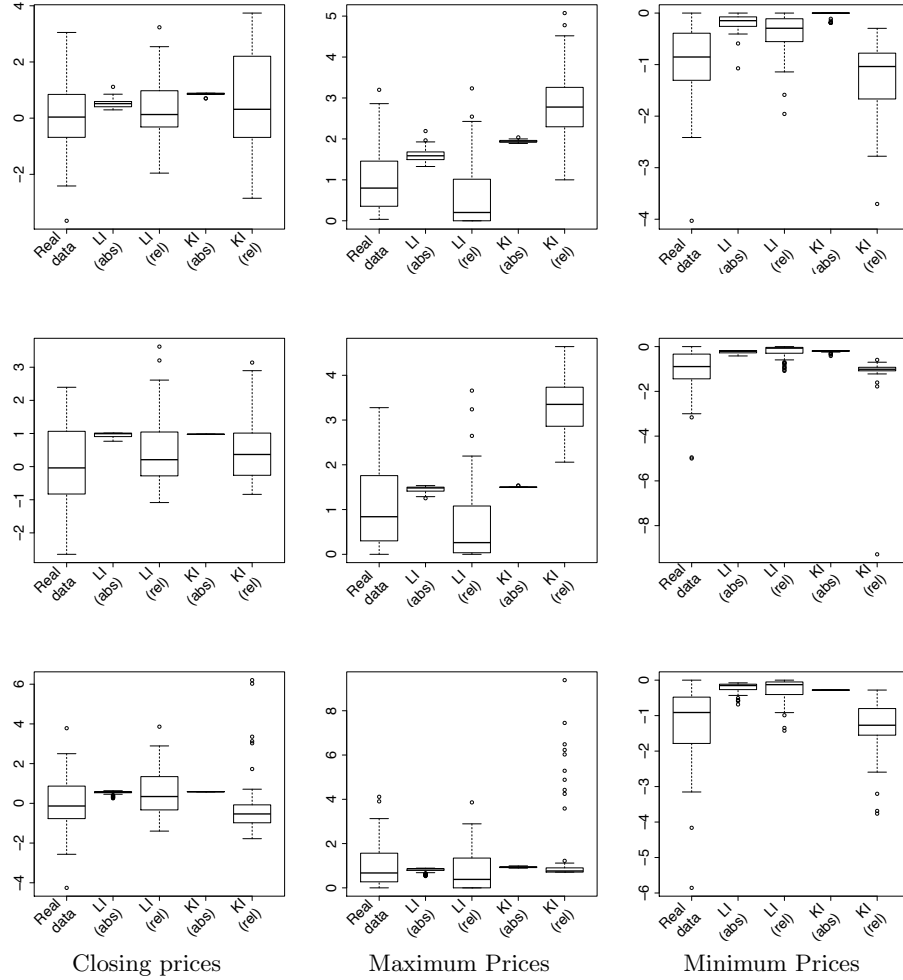


Fig. 4: Variation of the percentage difference between starting prices and closing, maximum and minimum prices observed in real data and in runs of our semi-synthetic model for GDF Suez(top), Arcelor Mittal(middle) and Deutsche Bank(bottom)

homogeneous construction of LI traders, and the end result is a time series that doesn't meet the volatility clustering requirement.

4.1 Answers to research questions

Regarding the effect of absolute pricing, our results indicate that using this assumption for the behaviour of traders in the real dataset produces a very narrow range of prices and thus cannot be used in a model that aims to create

alternative scenarios for the short-term behaviour of a stock price. Using relative pricing, on the other hand, allows the price in the simulated market to drift away from the price in the real market and thus the model can create a reasonable range of prices.

Regarding the effect of simple (LI) traders, their homogeneity is problematic, as although batch runs (with relative pricing) produce reasonable ranges of prices, we also see unrealistic stock price behaviour, with constant upward or downward trends. In addition, the time series of stock prices in many simulation runs fails the volatility clustering requirement that we need to show similarity to real financial time series. Using realistic (KI) traders with relative pricing produces reasonable ranges again, but also meets the volatility clustering requirement. While the price ranges do not match up exactly with the ranges produced from the real dataset (particularly with regards to the maximum prices observed in the day), our results demonstrate that of the four combinations tested, relative pricing with KI traders produces simulation runs that are closest to real financial time series overall, taking into account the range of prices produced in the simulation, as well as the presence of fat tails in the returns distribution and of volatility clustering in the prices.

5 Limitations and threats to validity

The categories of traders mentioned in Kirilenko et al. [9] provide only a high level description of the strategies followed. As such, we have had to select what we believe are sensible values for a number of parameters. Future work will repeat the tests to measure the sensitivity of our results to small changes in the parameters.

In addition, we have only tested the semi-synthetic model on the intra-day stock prices of 3 companies, which limits the degree to which we can generalise our results. Future work will extend these results to more companies to provide a more complete picture of the behaviour of our agent-based model.

6 Related work

Research regarding the simulation of financial markets using agent-based modelling can be traced back 30 years ago to the work of Cohen et al. [5], who proposed a model for a stock exchange. Cohen evaluated the impact of various stabilising policies on price, volatility and liquidity. He also introduced the concept of heterogeneous trading agents and an architecture for the limit order book, ideas which have been replicated in many forms since.

More recently, a variety of approaches have been suggested, each drawing from a wide and varied literature, including Finance, Economics, Mathematics, Statistics and Physics. The aims of these approaches vary, from trying to implement ‘rational’ models of trader behaviour to reproducing particular statistical features of markets. More details can be found in the surveys of LeBaron [10] and Chakraborti et al. [2]. Cristelli et al. [6] studied the commonalities, strengths

and deficiencies of existing models and proposed additional questions to be considered in future models.

The Santa Fe Artificial Stock Market ([16, 1, 11]), is one of the best-known examples of agent-based financial markets. Santa Fe agents make trading decisions based on binary market descriptors, and their strategies evolve in order to maximize profitability. The papers above also deal with the rate of evolution of the strategies, and how this gives rise to different regimes; the rational expectations regime and a more complex regime, where bubbles and crashes may appear. Other evolutionary approaches can be found in Chen and Yeh [3], Lux and Schornstein [13] and Pereira et al. [17].

Another strand of research in the agent-based modelling of financial markets literature attempts to recreate market characteristics by giving agents more realistic strategies, similar to the second class of traders implemented in this paper. This was initially attempted by Kim and Markowitz [8] and has been studied more recently by Westerhoff and Reitz [19] and Chiarella and Iori [4].

7 Conclusions

In this paper we have introduced semi-synthetic agent-based modelling, a new concept in the area of agent-based modelling of financial markets. Intuitively, a model for a stock market that is partially based on real order book data and partially based on agents should be closer to the real market than a pure agent-based model, in terms of simulated price dynamics. Our tests generate realistic runs of daily trading, when assuming that traders from the real markets (i.e. the traders that had submitted the orders in the real dataset) would have submitted their orders to buy or sell stocks not as suggested in the dataset, but shifted by the difference between the price they see in the simulated market and the price they had seen on the real market. If, in contrast, traders would have submitted their orders at exactly the same price, simulations that use both simple and realistic trading agents to interact with the real orders yield price dynamics which closely mimic those observed during the actual trading day, which is not useful when seeking to investigate alternative possible scenarios.

We investigated the effect of using homogeneous, limited intelligence traders compared to more varied, realistic traders, for the synthetic part of our model. Using realistic traders gave us the closest match with real markets, in terms of the market characteristics we measured. Using almost random traders gave us good results with regards to the range of prices achieved, but studying individual runs showed that the homogeneity of these traders gives rise to price behaviour not normally associated with intra-day price dynamics (constant upward or downward trends, as well as atypical volatility).

We believe that these results are promising, as they show that from a limited dataset (one trading day, in particular), we can generate thousands of realistic alternative scenarios. We hope that by extending this research, it will be possible to identify potential problems, like intra-day booms and crashes, and consider the impact from a range of policy measures.

References

1. Arthur, W.B., Holland, J.H., LeBaron, B.D., Palmer, R.G., Tayler, P.: Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. SSRN eLibrary (1996)
2. Chakraborti, A., Muni Toke, I., Patriarca, M., Abergel, F.: Econophysics: Empirical facts and agent-based models. ArXiv e-prints (Sep 2009)
3. Chen, S.H., Yeh, C.H.: Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market. *Journal of Economic Dynamics and Control* 25(3-4), 363–393 (Mar 2001)
4. Chiarella, C., Iori, G.: A simulation analysis of the microstructure of double auction markets. *Quantitative Finance* 2(5), 346–353 (2002)
5. Cohen, K.J., Maier, S.F., Schwartz, R.A., Whitcomb, D.K.: A simulation model of stock exchange trading. *Simulation* 41(5), 181–191 (1983)
6. Cristelli, M., Pietronero, L., Zaccaria, A.: Critical Overview of Agent-Based Models for Economics. ArXiv e-prints (Jan 2011)
7. Engle, R.F.: A general approach to lagrange multiplier model diagnostics. *Journal of Econometrics* 20(1), 83–104 (Oct 1982)
8. Kim, G., Markowitz, H.: Investment Rules, Margin, And Market Volatility. *Journal of Portfolio Management* 16(1), 45–52 (1989)
9. Kirilenko, A.A., Kyle, A.P., Samadi, M., Tuzun, T.: The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. Social Science Research Network Working Paper Series (Oct 2010), <http://ssrn.com/abstract=1686004>
10. LeBaron, B.: Agent-based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control* 24(5-7), 679–702 (2000)
11. LeBaron, B., Arthur, W., Palmer, R.: Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23(9-10), 1487–1516 (1999)
12. Licalzi, M., Pellizzari, P.: Fundamentalists clashing over the book: a study of order-driven stock markets. *Quantitative Finance* 3(6), 470–480 (2003)
13. Lux, T., Schornstein, S.: Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics* 41(1-2), 169–196 (Feb 2005)
14. Mandelbrot, B.: The Variation of Certain Speculative Prices. *The Journal of Business* 36(4), 394–419 (1963)
15. Maslov, S.: Simple model of a limit order-driven market. *Physica A: Statistical Mechanics and its Applications* 278(3-4), 571–578 (2000)
16. Palmer, R., Arthur, W.B., Holland, J.H., LeBaron, B., Tayler, P.: Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena* 75(1-3), 264–274 (1994)
17. Pereira, C.d.C., Mauri, A., Tettamanzi, A.G.B.: Cognitive-agent-based modeling of a financial market. In: Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology - Volume 02. pp. 20–27. WI-IAT '09, IEEE Computer Society, Washington, DC, USA (2009)
18. Russell, J.R., Engle, R.F.: Analysis of high-frequency data. In: Ait-Sahalia, Y., Hansen, L.P. (eds.) *Handbook of Financial Econometrics*, pp. 383–426. North-Holland, San Diego (2010)
19. Westerhoff, F., Reitz, S.: Nonlinearities and cyclical behavior: The role of chartists and fundamentalists. *Studies in Nonlinear Dynamics & Econometrics* 7(4), 3 (2003)
20. Zigrand, J.P., Shin, H.S., Beunza, D.: Feedback effects and changes in the diversity of trading strategies. *The future of Computer Trading in Financial Markets* (2011), <http://www.bis.gov.uk/assets/bispartners/foresight/docs/computer-trading/11-1221-dr2-feedback-effects-and-changes-in-diversity-of-trading-strategies.pdf>